

SVC-L2 VPNs: FLEXIBLE ON DEMAND SWITCHED MPLS/IP LAYER-2 VPNs FOR ETHERNET SVC, ATM AND FRAME RELAY

RELATED U.S. APPLICATION DATA

[0001] Provisional application No. 60/409,324 filed on September 9, 2002.

FIELD OF THE INVENTION

[0002] The present invention relates to switched virtual circuit (SVC) virtual private networks (VPNs) and is particularly concerned with flexible, on-demand switched MPLS/IP Layer-2 VPNs for Ethernet, ATM and Frame Relay SVCs.

BACKGROUND OF THE INVENTION

[0003] A Virtual Private Network (VPN) may be thought of as a private network constructed within a shared network infrastructure. In common terminology, these private networks are used by clients while the network infrastructure is supplied by providers.

[0004] Existing varieties of switched Layer-2 VPNs have limitations affecting ease of implementation and use including:

[0005] - clients must store and manipulate provider addresses;

[0006] - clients need to be configured with all the provider addresses to which the client has a site attached;

[0007] - clients need to know about connection restrictions, such as for closed-user-group (CUG) values, and need to signal these values when establishing connectivity;

[0008] - clients encounter complexity in managing CUG rules; and

[0009] - clients need to implement an appropriate Layer-2 signalling mechanism proper to the transport technology.

[0010] In view of the foregoing, it would be desirable to provide a technique for providing switched virtual circuit virtual private networks (SVC VPNs) which overcomes the above-described inadequacies and shortcomings.

SUMMARY OF THE INVENTION

[0011] An object of the present invention is to provide an improved switched virtual circuit Layer-2 virtual private network arrangement.

[0012] According to an aspect of the present invention, there is provided a network for providing switched virtual circuit Layer-2 VPNs, wherein the network includes a set of elements interconnected by services; at least one first subset of said elements defining a private network; and at least one second subset of elements different from said first subset defining a provider network wherein at least two subgroups of said first subset of elements may be connected via said provider network. There are a plurality of customer ports maintained on the elements of the first subset of elements and a plurality of provider ports maintained on the second set of elements, each of the plurality of provider ports connected by data and signalling services to a customer port. At each element of the provider network having a provider port is a port information table containing mapping information relating addresses of customer ports to addresses of provider ports for the first subset of elements. The network also includes a provisioning mechanism used to define element membership in said first subset of elements and a signalling mechanism used to create Layer-2 connectivity between elements within said first subset of elements at the Layer-2 level across said second subset of elements.

[0013] Advantages of the present invention include constrained and/or restricted connectivity defined by the customer but maintained and enforced by the provider. As well, the present invention provides for on-demand Layer-2 circuit requests. The demands initiate with the SVC-L2VPN customer and require no coordination with the provider in the usual sense of adding connections. The client devices operate within the SVC-L2VPN space independently from the provider network

operations in the sense of the provider network operations being transparent to the customer devices, yet connectivity is managed by the provider network relieving the customer of managing closed user groups. The present invention provides privacy and independence with respect to addressing, supporting an addressing unique to each SVC-L2VPN. The invention has the advantage of providing for single-ended provisioning, yet supports a multiservice Layer-2 switched model including Frame Relay, ATM, Ethernet, and Ethernet VLAN.

[0014] Conveniently the invention further provides for an auto-discovery mechanism for distributing said mapping information to port information tables of the provider network. This auto-discovery mechanism for distributing said mapping information uses Border Gateway Protocol in some instances.

[0015] In accordance with another aspect of the present invention, there is provided a method of organizing a network having a set of elements interconnected by services, wherein at least one first subset of the elements defines a private network and at least one second subset of elements different from the first subset defines a provider network and wherein at least two subgroups of the first subset of elements may be connected via the provider network, wherein the method includes the steps of defining element membership in the first subset of elements via a provisioning mechanism, establishing a plurality of customer ports within the elements of the first subset of elements and establishing a plurality of provider ports within the second set of elements. Each of the plurality of provider ports are connected by data and signalling services to a customer port. Thereafter, the step of establishing a port information table at each element of said provider network having a provider port, the port information table containing mapping information relating addresses of customer ports to addresses of provider ports, and creating Layer-2 connectivity within said first subset of elements at the Layer-2 level across said second subset of elements via a signalling mechanism.

[0016] The present invention further includes a method of organizing a network having a set of elements interconnected by services, wherein at least one first subset of the elements defines a private network, at least one second subset of elements different

from the first subset defines a provider network and wherein at least two subgroups of the first subset of elements may be connected via said provider network. The method includes the steps of defining a L2VPN topology; establishing a plurality of customer ports within said elements of the first subset of elements; and establishing a plurality of provider ports within the second set of elements, each of the plurality of provider ports connected by data and signalling services to a customer port. Thereafter, the steps of creating a Layer-2 Port Information Table for each provider port; establishing the identity of customer ports attached to each provider port, and populating the Layer-2 Port Information Table at that provider port with mapping information relating addresses of customer ports to addresses of provider ports. Subsequently, the mapping information is distributed to the Layer-2 Port Information tables of the provider network via an auto-discovery mechanism. The method then creates Layer-2 connectivity within the first subset of elements at the Layer-2 level across the second subset of elements via a signalling mechanism upon request from an element within the first subset of elements.

[0017] The present invention will now be described in more detail with reference to exemplary embodiments thereof as shown in the appended drawings. While the present invention is described below with reference to the preferred embodiments, it should be understood that the present invention is not limited thereto. Those of ordinary skill in the art having access to the teachings herein will recognize additional implementations, modifications, and embodiments which are within the scope of the present invention as disclosed and claimed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] The invention will be further understood from the following detailed description of embodiments of the invention and accompanying drawings in which:

[0019] **FIG. 1** is a diagram of a generic network having a shared network infrastructure and Virtual Private Networks associated thereto;

[0020] **FIG. 2** is a diagram of a network reference model including a plurality of customer edge devices, provider edge devices, and provider devices within the network;

[0021] **FIG. 3** is a diagram of the relation between Customer Ports and Provider Ports according to an embodiment of the invention;

[0022] **FIG. 4** is a diagram of a network depicting network addresses and Port Information Tables according to an embodiment of the invention;

[0023] **FIG. 5** is a diagram of a network depicting network addresses and Port Information Tables according to an alternative embodiment of the invention;

[0024] **FIG. 6** is a diagram of a network depicting network addresses according to an embodiment of the invention;

[0025] **FIG. 7** is a set of Layer-2 Port Information Tables depicting addressing updates via a BGP mechanism according to an alternative embodiment of the invention;

[0026] **FIG. 8** is a table depicting an example BGP Update Message according to an embodiment of the invention;

[0027] **FIG. 9** is a table depicting an example SVC-L2VPN Network Layer Reach Information tuple according to an embodiment of the invention.

DETAILED DESCRIPTION

[0028] Glossary of Acronyms Used

P – Provider Device

PE – Provider Edge Device

CE – Customer Edge Device

SVC – Switched Virtual Circuit

CPI – Customer Port Identifier (Layer-2)

PPI – Provider Port Identifier (Layer-2)

PIT – Port Information Table

L2PIT – Layer-2 Port Information Table

BGP – Border Gateway Protocol

BGP-AD – BGP Auto-Discovery

MPLS – Multi-Protocol Label Switching

DLCI – Data Link Connection Identifier

LMP – Link Management Protocol

ISP – Internet Service Provider

[0029] Referring to **FIG. 1**, there may be seen a generic network having a shared network infrastructure **100** with connected virtual private network sites **101**. The VPN sites **101** make use of the network infrastructure **100** to interconnect physically remote sub-networks of particular VPNs.

[0030] Referring to **FIG. 2**, there may be seen a network reference model showing a more detailed depiction of a network having a plurality of customer edge router/switches (CEs) **201, 202, 203, 204, 205, 206, 207, 208** and **209**. The provider network has provider edge router/Layer-2 switches (PEs) **210, 212**, and **214** as well as provider devices (P) **215, 216, 217**, and **218** interior to the provider network.

[0031] Further in **FIG. 2** may be seen the typical case where VPN A has a portion connected to CEs **201** and **202**, and another portion connected to CE **206**.

Communication services between these remote portions of VPN A will be provided by the provider network. The same general situation obtains for VPN B, VPN C, and VPN D.

[0032] In operation, the Switched Virtual Connection Layer-2 VPN (SVC-L2VPN) is a provider-based Layer-2 VPN service that allows clients to request on-demand Layer-2 circuits characterized by:

[0033] – a given topology;

[0034] – using IP/MPLS based signalling between CE-PE;

[0035] – the possible employment of Link management protocol (LMP) for Layer-2 link-port consistency;

[0036] – use of private addresses which have the potential to be overlapping with other addresses in other VPNs; and

[0037] – the capacity to be built using single-sided signalling and auto-discovery mechanisms as, for example, being standardized in IETF.

[0038] In one contemplated embodiment, IP/MPLS is used between CE and PE to convey signalling information, and between PEs (for example when the provider network is using IP/MPLS-based tunnelling).

[0039] A formulaic description would be as follows:

SVC MPLS/IP L2VPN \equiv SVC + (G)MPLS + IP + VPN Constructs

[0040] where:

SVC implements the private switched model;

(G)MPLS provides signalling for Layer-2 connections;

IP is the IP control channel; and

VPN Constructs are services such as VPN membership, overlapping addresses, VPN auto-discovery, etc.

[0041] Under SVC-L2VPNs, the following capabilities are provided:

[0042] **Constrained/restricted connectivity:**

- as defined by customer; and

- as maintained/enforced by the service provider.

[0043] On-demand L2 circuit request:

- controlled by the SVC-L2VPN customer;
- able to require no co-ordination with the service provider;
- client devices operate within the SVC-L2VPN space independently from the service provider network operations; and
- able to be subject to constrained/restricted connectivity.

[0044] Privacy/independence with respect to addressing within VPN.

[0045] Single ended provisioning.

[0046] Multi-Serviced Layer-2 Switched model:

- including for example, ATM, Frame Relay, Ethernet, and Ethernet VLAN (PPP, HDLC, etc.).

[0047] The use of an SVC-L2VPN allows for simplified provisioning. In the addition of a new port to a given SVC-L2VPN, the configuration and provisioning changes only on the PE that has this port. Typically BGP would be used to distribute the information to other PEs having ports of the given SVC-L2VPN. Likewise, BGP would also be used to distribute this information to other CEs that have ports of the given SVC-L2VPN. In terms of establishing/terminating a Layer-2 connection between a pair of ports in a given SVC-L2VPN, the client could execute the connection without involving configuration/provisioning changes in any of the Provider equipment by using (G)MPLS signalling.

[0048] A number of benefits for both client and provider are associated with SVC-L2VPNs as compared to legacy Layer-2 VPNs.

[0049] Advantages to the VPN Customer on the client side include:

- [0050]** - compatibility with access clients that are 'MPLS/IP' signalling based;
- [0051]** - support for overlapping and/or private address space;
- [0052]** - support of Layer-3 addresses within the L2VPN (does not require transport Layer-2 addresses);
- [0053]** - provides High Mobility capability in that the client can move its L2VPN from one port to the other without changing the addressing of the L2VPN and without changing the L2VPN addressing, QoS, etc., thus offering a

greater flexibility for network operations such as ATM to Ethernet, for example;

- [0054] - L2VPN addresses can be used for client Layer-3 networking;
- [0055] - support for a range of security capabilities including the Layer-2 security;
- [0056] - support for a range of QoS capabilities that includes Layer-2 VPNs QoS
- [0057] - support for the SVC-L2VPN circuit to be used as a legacy Layer-2 circuit or as an MPLS LSP within the client network if needed; and
- [0058] - does not require the client to implement full MPLS but just signalling protocol at the edges.
- [0059] Advantages to the Service Provider on the provider side include:
- [0060] - opportunity for new revenue opportunities to the ISPs;
- [0061] - support for Dynamic Membership distribution to ease circuit configuration and distribution;
- [0062] - capable of interworking with existing legacy Layer-2 VPNs;
- [0063] - provides opportunity to maximize yield from network investment on legacy layer-2 and IP/MPLS based infrastructure;
- [0064] - leverages existing provider skill level in layer-2 VPNs;
- [0065] - avoids requirement for tunnelling (including MPLS) between PE-PE (only when MPLS is used in the core);
- [0066] - support for reusing (G)MPLS for link, port constructs ;
- [0067] - support for single-sided signalling; and
- [0068] - allows Provider network operations to be completely decoupled from the customer L2VPNs unlike the case for legacy switched L2VPNs.
- [0069] The SVC-L2VPN protocol requirements are as follows:
- [0070] - at the CE:
- [0071] - support for MPLS signalling, for example RSVP-TE with SVC-L2VPN extensions but not necessarily MPLS forwarding; and
- [0072] - IP-based control channel, for example, IP tunnelling.
- [0073] - at the PE:
- [0074] - IP based control channel;

- [0075] – MPLS signalling; and
- [0076] – optionally an auto-discovery mechanism.
- [0077] The SVC-L2VPN VPN Architecture Components may be summarized as follows:
- [0078] – Access is Layer-2
- [0079] – Require an IP-based control channel for signalling purposes for port control such as RSVP-TE or CR-LDP, and LMP can be used between CE-PE;
- [0080] – Layer-2 discovery mechanism using the Layer-2 discovery for the Layer-2 port information;
- [0081] – Membership is defined in the same way as existing Layer-2 VPNs wherein the route-target or more precisely a global unique identifier identifies the destination address;
- [0082] – Ports and links are logical constructs that uses (G)MPLS functions; and
- [0083] – Signalling is MPLS based (packet side only) between CE-PE.
- [0084] The SVC-L2VPN Building Blocks may be summarized as follows:
- [0085] - Customer and Provider Ports;
- [0086] - A Layer-2 Port Information Table (L2PIT) which maintains mapping between customer ports and provider ports (at the edges of the service provider network) provides local CEs with the information about other ports in the SVC-L2VPN, and is defined on a per SVC-L2VPN basis or for all the SVC-L23VPNs connected to PE;
- [0087] - a Layer-2 BGP based auto-discovery mechanism (BGP-AD) used to determine and distribute information related to customer and provider ports to the PEs, and to populate the L2PIT with this information; and
- [0088] - an MPLS-signalling mechanism (preferably GMPLS based although other mechanisms may be used) to create connectivity within the set of client devices that are part of the same VPN at the Layer-2 level.
- [0089] In the SVC-L2VPN, connectivity is constrained/restricted by the Provider. The Customer may select any SVC-L2VPN topology within a defined set where the set

is controlled by the customer. The set may include hub-and-spoke, full-mesh, or other topologies. The provider restricts the customer's L2VPN topology to only the one in the set defined by the customer.

[0090] In the SVC-L2VPN discovery occurs via a two-step mechanism of:

[0091] – User-Side Port information discovery which is concerned with discovery of remote Layer-2 ports; and

[0092] – Network-side VPN discovery which is concerned with discovery of customer and provider port information.

[0093] Note that BGP multiprotocol extensions may be used as well as other discovery mechanisms such as DNS.

[0094] As part of the SVC-L2VPN, Customer and Provider ports are logical constructs. The Ports represent grouping of physical resources with similar characteristics for the purpose of switched L2VPN service. A single port may multiplex several Layer-2 connections.

[0095] There is a flexible relationship between physical ports, ports, links, link-bundles, channels and sub-channels and SVC-L2VPNs. A single physical port or fiber can support multiple links (e.g. multiple wavelengths per fiber, where each wavelength would be in its own SVC-L2VPN). Multiple physical ports can be combined to form one logical port. Ports within an VPN need not have the same characteristics. For example, it is possible to link Ethernet to ATM, FR to ATM, FR to Ethernet, etc. This allows maximum interworking capabilities. Administrative ownership of ports is orthogonal to the SVC-L2VPN membership of these ports and, because of this, ports within an VPN could belong to the same (intranet) or different (extranet) organizations.

[0096] Referring to **FIG. 3**, there is depicted a sub-network at the edge of a provider network having a PE **301** which connects to a first CE **303** and a second CE **305**. The Customer Port connects a Customer Device to the Service Provider. It forms the basic unit of switched L2VPN membership. A given port can be in exactly one VPN; however, a given CE may of course have ports from different VPNs. A particular Customer Port **304** at the connection of CE **303** and PE **301** is indicated.

[0097] The Provider Port connects a PE to a Customer Device. It forms the basic unit of L23VPN membership. A given port can be in exactly one L2VPN; however, a given PE may have ports from different L2VPNs. Association with a particular L2VPN is established and maintained by the Provider's provisioning system, and is susceptible to being changed only by the Provider's provisioning system. A particular Provider Port **302** at the connection of PE **301** and CE **303** is indicated.

[0098] In the SVC-L2VPN, addressing and routing used by the Provider network offering the service is completely independent from the addressing and routing used by the SVC-L2VPN clients of that network. For the purpose of the SVC-L2VPN service, addressing and routing used by one SVC-L2VPN client need not be co-ordinated with any other SVC-L2VPN clients. Several choices of addresses can be used between CE-PE such as IPv4, IPv6, or NSAP (encoded in IPv6, as per RFC1888).

[0099] Port information contains addressing information and may also include information about the characteristics of the data link within that port.

[0100] Customer Ports are identified by the Customer Port Identifier (CPI). The CPI could be either an address or a tuple such as <CE address, CE port index>. The CPI is unique within a given SVC-L2VPN, but need not be unique across multiple SVC-L2VPNs. The optional additional information about characteristics of the data link within that port may include items such as type (Frame Relay, ATM), bandwidth, total unreserved bandwidth within the port, etc.

[0101] Provider Ports are identified by the Provider Port Identifier (PPI). The PPI could be either an address or a tuple such as <PE address, PE port index>. The PPI is unique within a Provider. When used with legacy L2VPN, the PPI can conveniently be NSAP, E.164, or X.121.

[0102] The addressing and routing used by the Provider network offering the service is independent from the addressing and routing used by the SVC-L2VPN clients of that network. For the purpose of the SVC-L2VPN service, addressing and routing used by one SVL-L2VPN client need not be coordinated with an other SVC-L2VPN clients. Several choices of addresses may be used between CE and PE including for example IPv4, IPv6, and NSAP.

[0103] The VPN-PPI identifies a Provider Port (attached to the CE) using an address taken from the customer network (i.e., unique only within that VPN).

[0104] In operation, an SVC-L2VPN service may be built by the following steps:

[0105] Define a set of L2VPN topologies;

[0106] 2) Create a L2PIT for L23VPN service;

[0107] 3) Connect Customer Ports to Provider Ports;

[0108] 4) Learn the identity of the attached customer ports at each PE;

[0109] 5) Populate the L2PIT with remote port information via an auto-discovery mechanism;

[0110] 6) Learn at each CE about port information related to other CEs which are members of the same L23VPN and part of a common L2VPN topology set; and

[0111] 7) Request, initiated from the CEs to the service provider, to establish a Layer-2 connection to other CEs within the same SVC-L2VPN.

[0112] In normal practice, the customer defines the set of potential SVC-L2VPN topologies. This is normally expressed in terms of a set of customer ports and potential connectivity among the ports within the set. A wide range of topologies is contemplated such as hub-and-spoke, full-mesh, or arbitrary. Once defined by the customer, the Provider is responsible for restricting the customer's SVC-L2VPN topology to members in the set defined by the customer.

[0113] To create an L2PIT for SVC-L2VPN, it is necessary to configure an L2PIT on the PE, associate a list of export/import communities to the L2PIT according to the set of SVC-L2VPN topology selected, and populate L2PIT with port information.

[0114] Each L2PIT on a PE is populated from two sources:

[0115] - the Customer-Port-to-Provider-Port mapping information for the Customer Ports connected to the attached PE; and

[0116] - the Customer-Port-to-Provider-Port mapping information for the Customer Ports connected to other PEs.

[0117] For the Customer-Port-to-Provider-Port mapping information for the Customer Ports connected to the attached PE, only the ports that belong to the SVC-L2VPN are associated with the L2PIT. The use is local to the PE information. This

information is distributed to other PEs having ports belonging to the SVC-L2VPN by exporting this information into BGP, and tagging it appropriately with a specific “export” community.

[0118] For the Customer-Port-to-Provider-Port mapping information for the Customer Ports connected to other PEs, only the ports belonging to the SVC-L2VPN associated with the L2PIT are associated. This information is carried by the auto-discovery mechanism using BGP (with multi-protocol extensions).

[0119] For populating the L2PIT via an auto-discovery mechanism, BGP Auto-Discovery (BGP-AD) allows automatic discovery of SVC-L2VPN members with their associated port information. BGP-AD is used to populate the L2PITs with other port information (for non-local PEs). BGP Multi-protocol extensions are used to carry port information. BGP route filtering (based on BGP Extended Communities) is used to restrict distribution of this information to only the PITs of that SVC-L2VPN. Only the BGP routes with communities matching the import communities are used to populate the L2PIT.

[0120] Referring to **FIG. 4**, example L2PITs for two different VPNs may be seen for the network with addresses as indicated. On **FIG. 4**, the dashed lines **401** represent connections established for VPN A and the solid lines **403** represent connections established for VPN B. L2PIT **405** shows the four address pairs and additional information for VPN A on PE 1. L2PIT **407** shows the three address pairs and additional information for VPN B on PE 1.

[0121] Referring to **FIG. 5**, there may be seen an example L2PIT for a network with legacy L2VPNs with addresses as indicated. On **FIG. 5**, L2PIT **501** shows the four address pairs and additional information for VPN A on PE 1.

[0122] Referring to **FIGURES 6 and 7**, there may be seen an example of how BGP may be used to update a set of L2PITs associated with PEs. In **FIG. 6** a network with example addresses as indicated may be seen. The connections for VPN B are indicated as solid links. In **FIG. 7** the three L2PITs associated with respective PE's may be seen. At point **701**, the L2PITs have received their initial values from the local connection.

[0123] Next, a BGP update from PE 1 to PE 2 provides the following information:

[0124] < CPI = 10.1.1.1, PPI = 16.1.1.1, FR, Community = B >. At point **702**, the updated information is reflected in the L2PITs.

[0125] Next, a BGP update from PE 1 to PE 3 provides the following information:

[0126] < CPI = 10.1.1.1, PPI = 16.1.1.1, FR, Community = B >. At point **703**, the updated information is reflected in the L2PITs.

[0127] Next, a BGP update from PE 3 to PE 1 provides the following information:

[0128] < CPI = 10.1.1.2, PPI = 16.1.1.5, FR, Community = B >. At point **704**, the updated information is reflected in the L2PITs.

[0129] Next, a BGP update from PE 2 to PE 1 provides the following information:

[0130] < CPI = 10.1.1.3, PPI = 16.1.1.7, FR, Community = B >. At point **705**, the updated information is reflected in the L2PITs.

[0131] Next, a BGP update from PE 3 to PE 2 provides the following information:

[0132] < CPI = 10.1.1.2, PPI = 16.1.1.5, FR, Community = B >. At point **706**, the updated information is reflected in the L2PITs.

[0133] Next, a BGP update from PE 2 to PE 3 provides the following information:

[0134] < CPI = 10.1.1.3, PPI = 16.1.1.7, FR, Community = B >. At point **707**, the updated information is reflected in the L2PITs.

[0135] At this point all the L2PITs are fully populated.

[0136] A CE with one of its ports in a given L2VPN is provided with the information about all other customer Layer-2 ports of that L2VPN. This is provided by the local (directly connected) PE as part of the SVC-L2VPN service. The information includes CPIs of the ports and may also include information about characteristics of the channels within these ports e.g., FR, ATM, bandwidth, etc. The information represents a subset of the L2PIT that is maintained by the PE for the SVC-L2VPN and does not include PPI information. The CE acquires this information from its local PE by using a variety of methods such as BGP, LMP, or other similar mechanisms. The CE uses this information to establish the actual layer-2 connectivity.

[0137] Establishing connectivity between CEs is controlled by CEs. A CE with a port in a given SVC-L2VPN is not required to have a connection to every other CE with

ports in that SVC-L2VPN. Connectivity is established when desired using the information that CE acquired from PE about other customer ports of that SVC-L2VPN.

[0138] A wide range of MPLS signalling protocols can be used including CR-LDP, RSVP-TE, and LDP-DOD although other signalling protocols are not excluded. The service provider uses their existing signalling protocols to establish the connection (e.g., Martini) once a request is initiated by a CE.

[0139] SVC-L2VPN Signalling comprises user-side signalling and network-side signalling. The user-side signalling is concerned with negotiating user-side Layer-2 parameters. The network-side signalling handles both legacy Layer-2 networks and IP/MPLS networks including both Martini VC signalling and single-sided signalling.

[0140] SVC-L2VPN can use concepts of single-sided signalling with some extensions. For an end to end connection, an SVC is established via bi-directional LSPs set-up between PEs representing the emulated VC. Each LSP is identified by, for example using terminology borrowed from single sided signalling, a tuple such as:

<{CPI,PPI},PE1, Attachment identifier VC1, PE2, Attachment Identifier VC2 >

[0141] The LSP in the opposite direction will be identified by a tuple such as:

<PE2, Attachment identifier VC2, PE1, Attachment Identifier VC1>

[0142] The required SVC is built by the pair of such LSPs.

[0143] When a signalling message is sent from PE1 to PE2, and PE1 needs to refer to an Attachment Identifier which has been configured on one of its own Attachment VCs (or pools), the Attachment Identifier is called a "Source Attachment Identifier" (SAI). If PE1 needs to refer to an Attachment Identifier which has been configured on one of PE2's Attachment VCs (or pools), the Attachment Identifier is called a "Target Attachment Identifier" (TAI). The Martini signalling will carry the source CPI,PPI and the target CPI,PPI. A PE which receives a Label Mapping Message containing a TAI will be able to map that TAI uniquely to one of its Attachment VCs (or pools). The way in which a PE maps a TAI to an Attachment VC (or pool) is as follows:

[0144] First, the source CPI/PPI will be matched to the list of entries into the L2VPN PIT if a match is found, an Attachment Identifier is created for the VC, and a PATH message is sent to the CE. Second, a label mapping message is also sent to the

remote endpoint with the source CPI, PPI taken from destination CPI,PPI of the first label mapping.

[0145] It is contemplated that SVC-L2VPN Extensions to MPLS Signalling may reuse some of the Martini draft with extensions of single sided signalling. For this to occur, then CE-PE extensions needs to carry the Type of Layer-2 circuits, needs to signal the type to the CE that the original CE may decide to initiate connectivity to, and needs to carry the CPI in the RSVP-TE/CR-LDP messages.

[0146] When Martini type signalling is used the following extensions are needed:

- the type of switched service;
- remote endpoint will need to signal the DLCI (data link connection identifier) to the CE;
- the CE may accept or terminate the VC; and
- the signalling needs to carry the <source CPI,PPI>,<target CPI,PPI>.

[0147] FIG. 8 provides an example BGP Update message for a SVC-L2VPN MP_REACH_NLRI attribute.

[0148] FIG. 9 provides an example SVC-L2VPN NLRI (Network Layer Reach Information) tuple. Optionally a Route Distinguisher (RD) may be encoded, but this is not strictly needed since PPI and CPI are normally unique to a given SVC L2VPN. If there is no PPI then an RD is necessary.

[0149] In the Frame Relay case, QoS requires Layer-2 Traffic Parameter negotiation. The following items needs to be signalled between CE-PE:

- [0150]** – Maximum outgoing and incoming frame information field size;
- [0151]** – Outgoing and incoming CIR;
- [0152]** – Minimum acceptable outgoing and incoming CIR;
- [0153]** – Outgoing and incoming committed burst size (Bc); and
- [0154]** – Outgoing and incoming excess burst size (Be).

[0155] In an alternative embodiment, Ethernet Point-to-Point SVCs may be established using the SVC-L2VPN arrangements discussed above. Since the virtual connection is accomplished through IP based mechanisms at the CE-PE, it is possible

to establish Ethernet-based SVCs in the same manner as above. Likewise, VLAN point-to-point SVCs can also be supported.

[0156] While the invention has been described in conjunction with specific embodiments thereof, it is evident that many alternatives, modifications, and variations will be apparent to those skilled in the art in light of the foregoing description. Accordingly, it is intended to embrace all modifications, variations and adaptations such as may be made to the particular embodiments of the invention described above without departing from the scope of the invention, which is defined in the claims.